The cognitive roots of regularization in language

Vanessa Ferdinand[a,b,*], Simon Kirby[b], Kenny Smith[b]

[a] *Santa Fe Institute, United States*
[b] *University of Edinburgh, United Kingdom*

arXiv:1703.03442v1 [cs.CL] 9 Mar 2017

Abstract

Regularization occurs when the output a learner produces is less variable than the linguistic data they observed. In an artificial language learning experiment, we show that there exist at least two independent sources of regularization bias in cognition: a domain-general source based on cognitive load and a domain-specific source triggered by linguistic stimuli. Both of these factors modulate how frequency information is encoded and produced, but only the production-side modulations result in regularization (i.e. cause learners to eliminate variation from the observed input). We formalize the definition of regularization as the reduction of entropy and find that entropy measures are better at identifying regularization behavior than frequency-based analyses. We also use a model of cultural transmission to extrapolate from our experimental data in order to predict the amount of regularization which would develop in each experimental condition if the artificial language was transmitted over several generations of learners. Here we find an interaction between cognitive load and linguistic domain, suggesting that the effect of cognitive constraints can become more complex when put into the context of cultural evolution: although learning biases certainly carry information about the course of language evolution, we should not expect a one-to-one correspondence between the micro-level processes that regularize linguistic datasets and the macro-level evolution of linguistic regularity.

Keywords: regularisation, frequency learning, domain generality, domain specificity, language evolution

1. Introduction

Languages evolve as they pass from one mind to another. Immersed in a world of infinite variation, our cognitive architecture constrains what we can perceive, process, and produce. Cognitive constraints, such as learning biases, shape languages as they evolve and also can help explain the structure of language (Kirby et al., 2014). Early on, debate over the nature of these biases was polarized: Chomsky's nativism program explained linguistic structure as the product of a language-specific acquisition device (Chomsky, 1957) while behaviorists claimed general-purpose learning mechanisms, such as reinforcement learning, could explain language acquisition (Skinner, 1957). Recent experimental research has found that domain-general learning mechanisms underwrite many aspects of language learning (Saffran & Thiessen, 2007), such as the statistical learning involved in word segmentation by infants (Saffran et al., 1996) and how memory constraints modulate learners' productions of probabilistic variation in language (Hudson Kam & Chang, 2009). However, it is likely that a mixture of domain-general and domain-specific mechanisms are involved in language learning (e.g. Pearl & Lidz, 2009; Culbertson & Kirby, 2015).

This paper offers a first attempt to quantify the relative contribution of domain-general and domain-specific learning mechanisms to linguistic regularization behavior. Regularization is a well-documented process by which learners impose structure on linguistic data by reducing the amount of variation in that data. When language learners encounter linguistic elements in free variation, such as two realizations of a particular phoneme, two synonyms for one meaning, or two possible word orders for constructing a clause, they tend to reduce that free variation by either eliminating one of the variants, or conditioning their variant use on some aspect of the context (e.g. on the adjacent linguistic context). Natural languages rarely exhibit free (i.e. unconditioned) variation (Givón, 1985) and the regularization behavior of language learners is likely to be the cause. Regularization has been documented extensively in natural language use and in the laboratory. In natural language, regularization occurs in children's acquisition of language (Berko, 1958; Marcus et al., 1992; Singleton & Newport, 2004; Smith et al., 2007), during the formation of creole languages from highly variable pidgin languages (Bickerton, 1981; Sankoff, 1979; DeGraff, 1999; Lumsden, 1999; Meyerhoff, 2000; Becker & Veenstra, 2003), during the formation of new signed languages (Senghas et al., 1997; Senghas, 2000; Senghas & Coppola, 2001), and in historical trends of language change (Schilling-Estes & Wolfram, 1994; Lieberman et al., 2007; van Trijp, 2013). In the laboratory, regularization has been

*Corresponding author at: Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, United States. Tel.: +1 505 946-2735.
Email address: vanessa@santafe.edu (Vanessa Ferdinand)

studied in depth through artificial language learning experiments with children (Hudson Kam & Newport, 2005, 2009; Wonnacott, 2011; Culbertson & Newport, 2015) and adults (Wonnacott & Newport, 2005; Reali & Griffiths, 2009; Smith & Wonnacott, 2010; Perfors, 2012; Culbertson et al., 2012; Fehér et al., 2016; Smith et al., 2017).

Behavioral experiments offer special insight into the regularization process, because they allow researchers to present participants with controlled linguistic variation, precisely measure the way participants transform that variation, and test hypotheses about what causes participants to alter patterns of variation. For example, Hudson Kam & Newport (2009) investigated the regularization of pseudo-determiners in an artificial language learning experiment. Participants were trained on a language that consisted of several verbs, several nouns (divided into 2 noun classes), 2 *main* determiners (one for each noun class), and zero to 16 *noise* determiners (which could occur with any noun). In the training language, each noun occurred with its main determiner on 60% of exposures; the remaining exposures were equally divided across the noise determiners. In the testing phase, participants described scenes using the language they had learned. When participants encountered only two noise determiners during training, they regularized slightly by producing the main determiners with 70% of the nouns, rather than the 60% they observed in the training language. Regularization increased with the number of noise determiners, reaching its highest level with 16 noise determiners, where the main determiners were produced with nearly 90% of the nouns. A followup experiment in Hudson Kam & Newport (2009) shows that participants regularize the same artificial language less when the noise determiners are conditioned on particular nouns in a more predictable and consistent way.

These results are consistent with Newport's Less-is-More hypothesis. Originally conceived as an explanation for why children regularize more than adults (Newport, 1990), it states that learners with limited memory capacity may regularize inconsistent input because they have more difficulty storing and retrieving forms that are lower in frequency or used less consistently. This constitutes a domain-general account of linguistic regularization in terms of cognitive constraints on memory encoding and retrieval. If the Less-is-More hypothesis describes a truly domain-general effect, we should expect to see the same kind of regularization behavior in non-linguistic domains.

Gardner (1957) conducted a frequency prediction experiment in which adult participants had to predict which of several lights would flash in any given trial. When participants observed two lights flashing at random in a 60:40 ratio (light A flashed 60% of the time and light B flashed 40% of the time), they *probability matched* this ratio in their predictions, meaning that about 60% of their guesses were that light A would flash next and about 40% of their guesses were on light B. They also probability matched when observing a 70:30 ratio. However, when participants were trained on three lights (four ratios were tested: 70:15:15, 70:20:10, 60:20:20, and 60:30:10), they regularized by over-predicting the most frequent light and under-predicting the less frequent lights, which is similar to the behavior of Hudson Kam & Newport's (2009) participants. In another experiment, Kareev et al. (1997) report an effect of individual differences in working memory capacity (as determined by a digit-span test) on participants' perception of the correlation of two probabilistic variables. Participants with lower memory capacity overestimated the most common variant, whereas participants with higher capacity did not. Similarly, Dougherty & Hunter (2003) show that participants with lower working memory are less likely to consider alternative choices in an eight-item prediction task and are less likely to consider the low-frequency alternatives than participants with higher working memory. Each of these cases can be identified as regularization where the higher-frequency variants are over-represented in participants' behavior.

There is therefore strong evidence for the existence of domain-general drivers of regularization, but the extent to which they account for the level of regularity that we observe in language is not clear. This is because domain-specific learning mechanisms may play a role on their own, or interact with general mechanisms. For example, Perfors (2012) presents seven carefully controlled manipulations of cognitive load during the encoding stage of an artificial language learning task and finds no effect on regularization behavior. This suggests that the Less-is-More Hypothesis may apply more to retrieval than to storage, and that the effects of working memory found in the non-linguistic experiments of Kareev et al. (1997) and Dougherty & Hunter (2003) may not operate as strongly in language learning. Furthermore, Reali & Griffiths (2009) show an effect of domain on regularization behavior: participants reduce variation when learning about words but increase variation when learning about coin flips. However, cognitive load was lower in the coin flipping condition (one coin was flipped, whereas 6 objects were named), so it is unclear whether the higher cognitive load or linguistic domain caused participants to regularize in the word learning task.

In the following, we present a two-by-two experimental design that manipulates cognitive load (following Hudson Kam & Newport (2009) and Gardner (1957)) and task domain (directly comparing regularization in linguistic and non-linguistic domains). To manipulate cognitive load we vary the number of stimuli a learner must track concurrently. We manipulate task domain by manipulating the type of stimuli the learner must track: objects being named with words (linguistic domain) or marbles being drawn from containers (non-linguistic domain). Our method is closely based on the artificial language learning experiment in Reali & Griffiths (2009) and our high load linguistic condition replicates their Experiment 1.

Based on the work reviewed above, we predict that regularization behavior will increase when cognitive load is raised. We also predict that regularization behavior will

increase when the task is presented with linguistic stimuli. However, we have no clear prediction about the existence of an interaction between domain and cognitive load, or the relative amount of variation that will be removed from the data due to load or domain. Knowing the relative contribution of domain-general and domain-specific biases to structure in language is important because it tells us how much we can ground our theories of language learning in general mechanisms of memory and statistical learning. Furthermore, the extent to which language learning biases are domain-general tells us the extent to which the cognitive capacities which underpin language could have gradually developed from the cognition of our animal ancestors (e.g. Hauser et al., 2002; Fehér, 2016).

In order to address these questions, we need a principled measure of regularization that is comparable across different distributions of variation and stimuli domains. In the following section, we provide this measure by formalizing the definition of regularization as the reduction of entropy in a data set. Next we present our experimental method and results, where we compare how cognitive load and linguistic task domain elicit regularization behavior. Last, we use these measurements to explore the evolution of regularity as learners' biases are repeatedly applied under a model of cultural transmission. This will give us a sense of how predictive known regularization biases can be for the level of regularity found in languages.

2. Defining and quantifying regularization

In the existing literature, regularization is described as the elimination or reduction of free variation. Therefore, we will define regularization in terms of this lost variation and quantify it as the amount of variation that was lost from learners' productions when compared to the data the learners observed. The amount of variation in any data set can be quantified by the information-theoretic notion of entropy (e.g. Cover & Thomas, 1991) and a growing number of studies are using entropy measures to analyze regularization behavior (e.g. Smith & Wonnacott, 2010; Perfors, 2012; Fedzechkina, 2014; Ferdinand, 2015; Cuskley et al., 2015; Perfors, 2016; Fehér et al., 2016; Smith et al., 2017).

The variation in a distribution of items, such as linguistic variants, can be quantified by *Shannon entropy* (Shannon, 1948):

$$H(V) = -\sum_{v_i \in V} p(v_i) \log_2 p(v_i) \quad (1)$$

where V is the set of linguistic variants in question, $p(V)$ is the probability distribution over those variants, and $p(v_i)$ is the probability of ith variant in that set. For example, take the probability distribution over the 4 determiners used in the "2 noise determiner" condition of Hudson Kam & Newport (2009)'s artificial language learning experiment: $p(V) = \{0.3, 0.3, 0.2, 0.2\}$. The Shannon entropy of this distribution is 1.97 bits. Imagine a participant who was trained on this language and on testing produced the distribution $p(V') = \{0.7, 0.1, 0.1, 0.1\}$. The Shannon entropy of $p(V')$ is 1.36 bits and the change in variation is -0.61 bits. This means that 0.61 bits of variation among determiners was regularized (i.e. removed) by the participant. Or, more intuitively, $\frac{0.61}{1.97} \cdot 100 = 31\%$ of the variation in determiners was regularized by the participant.

Variation can also be lost when variants become conditioned on other linguistic variables or contexts. For example, each determiner may have a conditional probability $p(v_i|c_j)$ of being produced with a particular noun class c_j, such that if one knows the class of the noun, one is better able to predict which determiner a speaker of that language will use with that noun. The variation in a distribution of items, after a conditioning variable is taken into account, is quantified by *conditional entropy* (Shannon, 1948):

$$H(V|C) = -\sum_{c_j \in C} p(c_j) \sum_{v_i \in V} p(v_i|c_j) \log_2 p(v_i|c_j) \quad (2)$$

where V is the set of linguistic variants and C is the set of conditioning contexts. Again, $p(V)$ is the probability distribution over variants, $p(C)$ is the probability distribution over contexts, $p(v_i|c_j)$ is the conditional probability of observing the ith variant in the jth context, and $p(c_j)$ is the probability that the jth context occurs. Given the format of this equation, we can see that the conditional entropy is the sum of the entropy of variants per context, weighted by the probability of each context. Assume for a moment that the $p(V)$ distribution over determiners is not conditioned on each noun class, meaning that all determiners have the same conditional probabilities regardless of the noun class they are used with, for example: $p(v_i|c_1) = \{0.3, 0.3, 0.2, 0.2\}$ and $p(v_i|c_2) = \{0.3, 0.3, 0.2, 0.2\}$. Assume also that any noun has the following probabilities of being in noun class 1 or 2: $p(C) = \{0.6, 0.4\}$. Let us call this mapping A. The conditional entropy of mapping A is 1.97 bits, identical to the entropy of the determiners themselves, because the noun class carries no information about which determiner is used. We can contrast this with another mapping, mapping B, where determiner use is conditioned on noun class such that $p(v_i|c_1) = \{0.5, 0.5, 0.0, 0.0\}$ and $p(v_i|c_2) = \{0.0, 0.0, 0.5, 0.5\}$. Here, the first two determiners in the set are exclusively produced with noun class 1 and the third and fourth determiners are exclusively produced with noun class 2. The conditional entropy of mapping B is 1.00 bit, while its entropy over determiners remains at 1.97 bits. If a participant had been trained on a language with mapping A and produced a language with mapping B, then they would have regularized 0.97 bits, or $\frac{0.97}{1.97} \cdot 100 = 49\%$ of the variation in mapping A.

By themselves, $H(V)$ and $H(V|C)$ do not fully describe the variation in a mapping between linguistic variants and contexts. This is because $H(C)$ can differ. The total

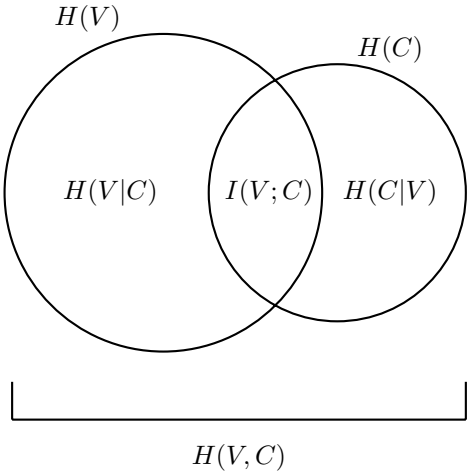

Figure 1: The relationship between entropy quantities in a mapping between linguistic variants (V) and their conditioning contexts (C).

amount of variation in a pair of random variables is given by the joint entropy:

$$H(V,C) = -\sum_{v_i \in V} \sum_{c_j \in C} p(v_i, c_j) \log_2 p(v_i, c_j) \qquad (3)$$

where $p(v_i, c_j)$ is the joint probability of observing the ith variant and the jth context together. For fixed values of $H(V)$ and $H(C)$, joint entropy increases as V and C carry less information about one another. In the case where V and C carry no information about one another (i.e. knowing what variant is being used tells you nothing about the context, and vice a versa), $H(V,C) = H(V) + H(C)$. This relationship allows us to calculate mutual information, which is how much uncertainty is reduced in V when C is known, $I(V;C) = H(V) - H(V|C)$ and how much uncertainty is reduced in C when V is known, $I(C;V) = H(C) - H(C|V)$. Note that $I(V;C) = I(C;V)$. Figure 1 shows the relationship between mutual information and five types of entropy that a mapping between linguistic variants and contexts can have.

Regularization is the reduction or elimination of entropy in a data set.

We define regularization as any reduction to the space in Figure 1. Regularization can occur by eliminating linguistic variants (reducing $H(V)$), eliminating conditioning contexts (reducing $H(C)$), or increasing the degree to which variants and contexts are conditioned on one another (reducing $H(V|C)$ and/or $H(C|V)$). Reduction in entropy is often accompanied by an increase in mutual information, however this is not always the case. Joint entropy always decreases when there is a net loss of variation. In the following experiment, we construct a stimuli set in which lexical items are variants and the objects they refer to are contexts. In a matched non-linguistic stimuli set, marbles are variants and the containers they are drawn from are contexts. The experiment is designed such that $H(V)$, $H(V|C)$, and $H(V,C)$ will always change by the same number of bits when participants regularize and $I(V;C)$ can not be changed.

3. Frequency learning experiment

In this experiment we manipulate cognitive load and task domain, allowing us to quantify the amount of variation participants regularize due to each source. Participants observe an input mapping among stimuli and then produce behavior from which an output mapping is extracted. Finally, they estimate the frequencies of the input stimuli and these estimates are compared to their output behavior.

3.1. Participants

573 participants were recruited via Amazon's Mechanical Turk crowdsourcing platform and completed our experiment online. Informed consent was obtained for experimentation. Participant location was restricted to the USA and verified by a post-hoc check of participant IP address location. 61 participants were excluded on the basis of the following criteria: failing an Ishihara color vision test (15), self-reporting the use of a pen or pencil during the task in an exit questionnaire (10), not reporting their sex or age (6), self-reporting an age below 18 (1), or having previously participated in this or any of our related experiments, as determined by their user ID with MTurk (26). More participants were recruited than necessary with the expectation that some would be excluded by these criteria. Once the predetermined number of participants per condition was met, the last participants were excluded (3). All participants (included and excluded) received the full monetary reward for participating in this task, which was 0.10 USD in the one-item conditions (*marbles1* and *words1*) and 0.60 USD in the six-item conditions (*marbles6* and *words6*). The average time taken to complete the one-item conditions was 3 minutes and 50 seconds, with a standard deviation of 1 minute and 27 seconds. The average time taken to complete the six-item conditions was 11 minutes and 32 seconds, with a standard deviation of 2 minutes and 6 seconds. Of the final 512 participants, 274 reported female, 238 reported male, and the mean age was 33.7 years (min = 18, max = 72) with a standard deviation of 11.3 years.

3.2. Materials and Stimuli

The experiment was coded up as a Java applet that ran in the participant's web browser in a 600x800-pixel field. Photographs of 6 different containers (a bucket, bowl, jar, basket, box, and pouch) and computer-generated images of marbles in 12 different colors (blue, orange, red, teal, pink, olive, lime, purple, black, yellow, grey, and brown) served as non-linguistic stimuli. Photographs of 6 different novel

objects (resembling mechanical gadgets) and images of 12 different nonsense words (*buv*, *kal*, *dap*, *mig*, *pon*, *fud*, *vit*, *lem*, *seb*, *nuk*, *gos*, *tef*) served as linguistic stimuli. Marbles and words were organized into fixed pairs that maximized distinctiveness between the stimuli in the pair. The stimuli lists above appear in order of these pairings (blue and orange were paired, buv and kal were paired, etc.). Marble colors were paired to differ in hue and brightness. Within-pair hue differences were greater than 120° (i.e. chosen from approximately opposite sides of the color wheel) and within-pair brightness differences were greater than 20%. Words were paired to be contrastive. Within-pair words utilized different letters and vowels and within-pair consonants differed by place of articulation. These stimuli are closely based on the word stimuli used in Reali & Griffiths (2009) and selected to not look or sound like existing words when pronounced by an American English speaker. Words were presented visually and were not accompanied by auditory stimuli.

3.3. Conditions and Design

We used a two-by-two design to investigate the effects of domain and cognitive load in four experimental conditions:

1) *Non-linguistic single frequency learning (marbles1)*
Participants observed two marble colors being drawn from one container at a particular ratio (for example, 5 blue marbles and 5 orange marbles displayed in random order). Participants were then asked to demonstrate what another several draws from the same container are likely to look like. They were not asked to predict specific future draws and thus no feedback was given. Participants observed 10 marble draws and produced 10 marble draws. Each participant observed a set of draws in one of six possible ratios: 5:5, 6:4, 7:3, 8:2, 9:1, and 10:0. These constitute six *input ratio conditions*. We will refer to the ratio that a participant observed as the *input ratio* and the ratio that the participant produced as the *output ratio*. There were 32 participants in each input ratio condition, totaling 192 participants in *marbles1*. Container stimuli were randomized across participants: each participant saw one of the six containers. Equal numbers of participants saw each container. Marble pairs were also randomized across participants: each participant saw one of the six marble pairs. Equal numbers of participants saw each marble pair. The full details of the observation and production regimes can be found in Section 3.4.

2) *Non-linguistic multiple frequency learning (marbles6)*
This condition is similar to the *marbles1* condition, with the difference that participants observed and produced 10 draws each from 6 different containers, where each container differed in the ratio of the two marble colors. Containers, marble pairs, and input ratios were randomly assigned to one another, without replacement, and these assignments were randomized between participants. Each participant saw all six of the containers, all six of the marble pairs, and all six of the input ratios (the same input ratios as were used in the *marbles1* condition: 5:5, 6:4, 7:3, 8:2, 9:1 and 10:0). There were 64 participants in this condition, yielding data for 384 (64x6) input ratios.

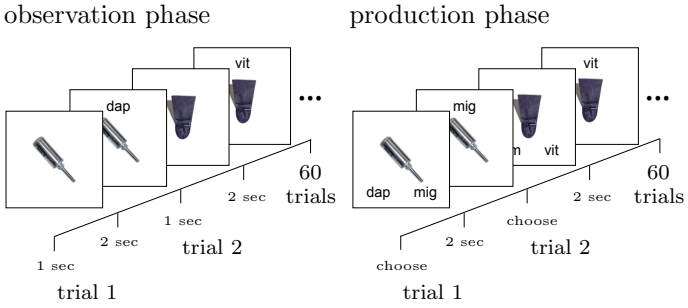

Figure 2: Schema of the observation and production phases of the experiment. Example shown is the high cognitive load linguistic condition. In the non-linguistic condition, containers are shown in place of the object, and marbles are shown in place of the words. In the low cognitive load conditions, participants only see 10 trials for the same object or container.

3) *Linguistic single frequency learning (words1)*
This condition is similar to the *marbles1* condition, differing only by the use of linguistic stimuli (objects and words) instead of the non-linguistic stimuli (containers and marbles) and minimal adaptation of the instructions to the linguistic domain. Participants observed one object being named with two words at a particular ratio (for example, *buv* 5 times and *kal* 5 times, in random order) and were then asked to name the object like they had observed it being named. They were not asked to predict specific future namings and thus no feedback was given. Participants observed 10 namings and produced 10 namings. Each of the 6 possible input ratios (same ratios as used in *marbles1*) was observed by 32 participants, totaling 192 participants for this condition.

4) *Linguistic multiple frequency learning (words6)*
This condition is similar to the *marbles6* condition, again differing only by the use of linguistic stimuli and minimal adaptation of the instructions to the linguistic domain. This condition constitutes a replication of the word learning experiment in Reali & Griffiths (2009), but with different object stimuli, modified word stimuli, and participants who completed the experiment online rather than in the laboratory. There were 64 participants in this condition, yielding data for 384 (64x6) input ratios.

3.4. Procedure

The experiment consisted of an observation phase and a production phase (Figure 2). In each observation trial, a container/object was displayed on its own for 1 second and then a marble/word was displayed above it for 2 seconds, with no break between trials. There were 10 observation

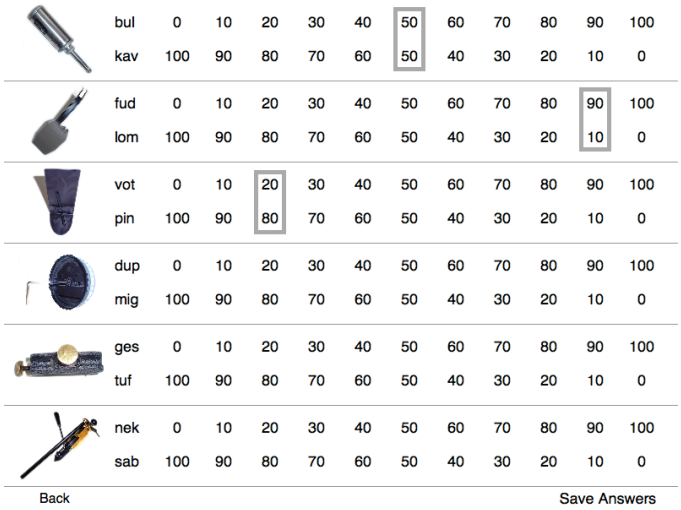

Figure 3: Screen shot of the sliders page in the high cognitive load linguistic condition, showing three answers selected. Participants could change their answers up until "Save Answers" was clicked. "Back" took participants back to the question and instruction about the sliders. In the low load condition, only one slider was shown.

trials per container/object. The majority marble/word in each pair was randomly assigned among the two stimuli. For example, in a 2:8 ratio, some participants would see the blue marble 8 times, whereas others would see the orange marble 8 times.

In each production trial, a container/object was displayed and the two marbles/words that appeared with it during observation were displayed below (one near the bottom left of the screen and the other near the bottom right). When participants clicked on one of the marbles/words, a small OK button appeared between the two choices. Participants could change their choice, but when they clicked on the OK button their current selection was registered as their answer and this choice was displayed above the container/object for 2 seconds. The OK button also served to center the participant's cursor between trials. There was no time limit on production trials. Production trials repeated until 10 responses were collected per container/object. Participants were not told how many observation or production trials there would be.

In the one-item conditions, participants received a total of 10 observation and 10 production trials. In the six-item conditions, there were a total of 60 observation and 60 production trials and the order in which the different containers/objects appeared was randomized across trials. For example, in observation trial 1 the participant might see one draw from the box, in trial 2 see one draw from the jar, trial 3 from the basket, trial 4 from the box again, and so on. Likewise, in production trial 1 the participant might be prompted to demonstrate a draw from the pouch, then from the basket in trial 2, and so on. The test-side locations of the two marbles/words were randomized per production trial (i.e. on each trial, each marble/word had an equal probability of being displayed on the left or right).

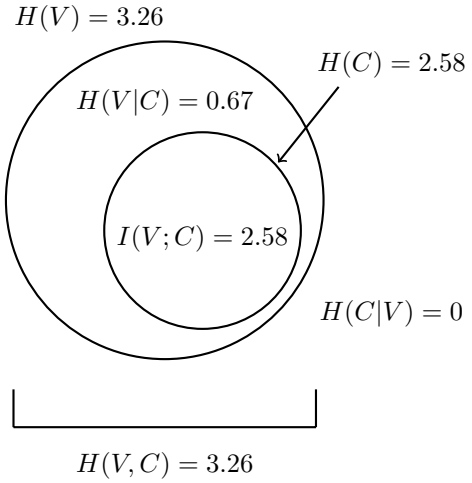

Figure 4: Entropy of the training stimuli (in bits). In the linguistic condition, V is the distribution over words and C is the distribution over objects. In the non-linguistic condition, V is the distribution over marbles and C is the distribution over containers. Refer back to Section 2 for the definition of each quantity. The experiment is designed so participants can change the size of the outer circle only.

	v_1	v_2	v_3	v_4	v_5	v_6	v_7	v_8	v_9	v_{10}	v_{11}	v_{12}
c_1	5	5	0	0	0	0	0	0	0	0	0	0
c_2	0	0	6	4	0	0	0	0	0	0	0	0
c_3	0	0	0	0	7	3	0	0	0	0	0	0
c_4	0	0	0	0	0	0	8	2	0	0	0	0
c_5	0	0	0	0	0	0	0	0	9	1	0	0
c_6	0	0	0	0	0	0	0	0	0	0	10	0

Table 1: Co-occurance frequencies among the twelve variants and six contexts in the experimental stimuli set. Each cell gives the number of times that the participant observed variant$_i$ along with context$_j$.

After the production phase, participants were asked to estimate the generating ratio that underlies the input ratio they saw. This was accomplished by asking participants how many marbles of each color were in each container, or how often each word is said for each object in the artificial language. Participants provided their response with a discrete slider over 11 options of relative percentages: 100:0, 90:10, 80:20, 70:30, 60:40, 50:50, 40:60, 30:70, 20:80, 10:90, 0:100 (Figure 3).

3.5. Entropy of the training stimuli set

Each participant observes a stimuli set that is composed of co-occurrances between marbles and containers or words and objects. For the purpose of quantifying the variation in the stimuli sets, we consider the marbles and words to be variants and consider the containers and objects to be contexts. Table 1 shows the co-occurrance frequencies between contexts and variants. In the high cognitive load conditions, this table describes the complete stimuli set that each participant was trained on in the observation phase. In the low cognitive load conditions, each participant was trained on only one row from the this table. Figure 4 shows the entropy values associated with Table 1 and describes the population-level variation

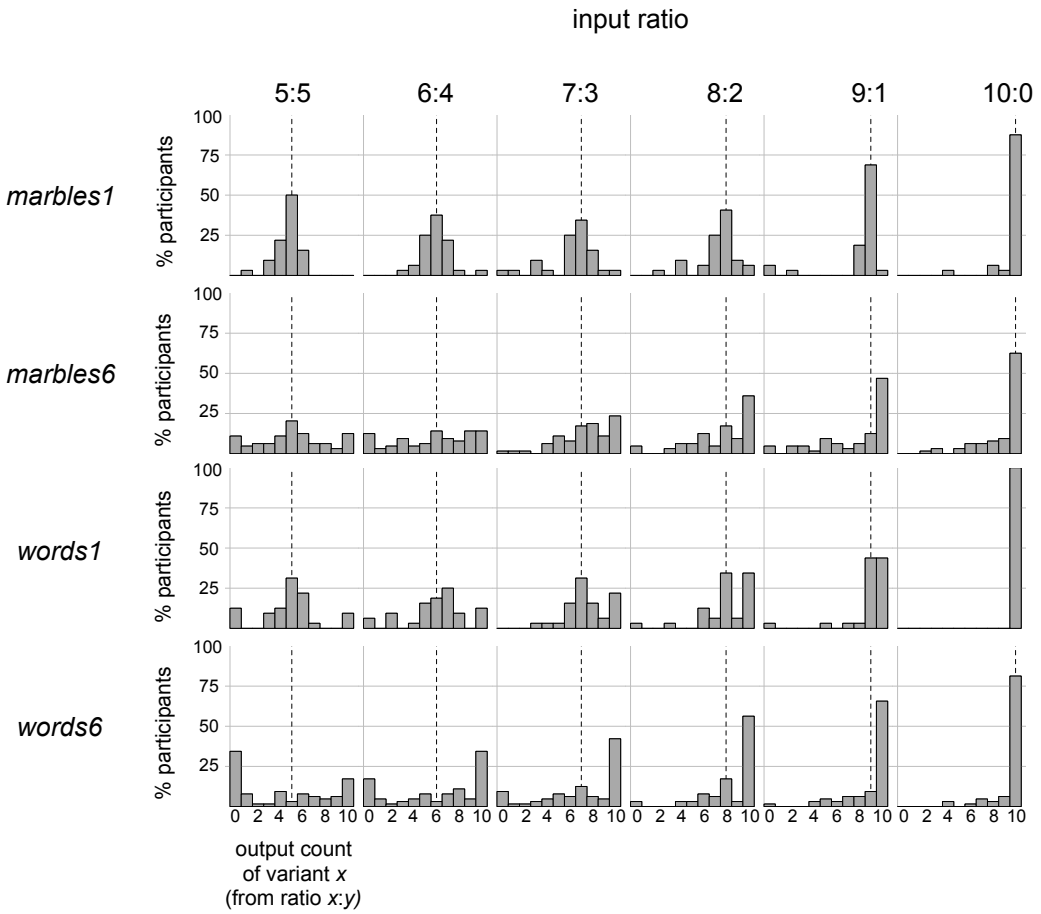

Figure 5: Each row shows the results of one experimental condition. Each column corresponds to one of the six input ratios, ranging from 5:5 (left) to 10:0 (right). Each pane contains the distribution of output ratios that participants produced in response to one input ratio. Output ratios are displayed on the x-axis as the number of times a participant produced variant x from the input ratio x:y, where variant x corresponds to whatever marble/word was in the majority during the observation phase. (In the 5:5 input ratio a random marble/word was coded as variant x.) All input ratios are indicated by a dashed line.

in stimuli. These values are the same across conditions, allowing the direct comparison of mean change in entropy between conditions.

It is important to note that the design of the production phase prevents participants from changing $H(C)$ because contexts are presented the same amount of times in the observation and production phases. The design also prevents participants from changing $H(C|V)$ because the only production options are the two variants that were shown with the context in the observation phase. Therefore, participants can only change the size of the larger circle in Figure 4. If participants regularize, $H(V)$, $H(V|C)$, and $H(V, C)$ will drop by the same number of bits.

The entropy of the stimuli that one participant observes in the high cognitive load condition is identical to Figure 4. However, the entropy of stimuli in the low cognitive load condition is lower and varies by the input ratio observed: in each input ratio condition, 5:5, 6:4, 7:3, 8:2, 9:1, and 10:0, $H(V) = H(V|C) = H(V, C) = I(V; C) =$ 1 bit, 0.97 bits, 0.88 bits, 0.72 bits, 0.47 bits, and 0 bits, respectively, and $H(C) = H(C|V) = 0$.

4. Results

4.1. Regularization behavior profiles

Before analyzing the data in terms of its entropy, we first visually inspect how participants changed each input ratio. In Figure 5, each panel shows the distribution of ratios that participants produced in response to each input ratio they observed, per experimental condition.

The first row (*marbles1*) shows clear probability matching behavior, where both the mean and mode of participant responses are near the input ratio. Participants in this condition tended to successfully reproduce their input ratio, with a small amount of error. The second row (*marbles6*) shows clear regularization behavior. Participants in this condition have moved distributional mass away from the input ratio and toward the maximally regular ratios, 0:10 and 10:0. Responses to the 5:5 input ratio seem to be a combination of probability matching behavior (13 participants also produced a 5:5 ratio) and regularization behavior (15 participants produced maximally regular ratios). The third row (*words1*) shows a mixture of probability

matching and regularization behavior for all input ratios. Roughly half of the participants appear to have probability matched with error rates similar to *marbles1*, and roughly half of the participants appear to have regularized at levels comparable to *marbles6*. In the 10:0 input condition, none of the participants choose the unseen word on any production trial. The fourth row (*words6*) shows a similar regularization profile to *marbles6*, but with a more extreme movement of distributional mass to the edges, such that the majority of participants produced maximally regular ratios. This condition constitutes a successful replication of the first experiment reported in Reali & Griffiths (2009).

4.2. Regularization per condition

In this section, we report the differences in regularization behavior within and between the four experimental conditions. We do this by calculating the change in Shannon entropy for each pair of input-output ratios obtained from participants. For example, if a participant observes a 5:5 ratio of orange and blue marbles for the jar, and then produces a 6:4 ratio of orange and blue marbles for the jar, the Shannon entropy for that pair of input-output ratios changes by -0.12 bits.[1] Figure 6 shows the mean change in entropy for all input-output ratio pairs per condition. Negative values mean participants made ratios more regular on average.

To assess the significance of differences in regularization within and between conditions, a linear mixed effects regression analysis was performed using R (R Core Team, 2013) and *lme4* (Bates et al., 2013). The dependent variable was the change in entropy of the input-output ratios. Experimental condition was the independent variable. Participant was entered as a random effect (with random intercepts). No obvious deviations from normality or homoscedasticity were apparent in the residual plots.

Within-condition changes were assessed by re-leveling the model to obtain the intercept value for each condition. The intercept equals the condition's mean change in entropy and the regression analysis provides a t-statistic to evaluate whether or not this mean is significantly different from zero. Three of the four experimental conditions elicited a significant amount of regularization behavior (Figure 6). Participants regularized an average of 0.17 bits in *marbles6*, 0.19 bits in *words1*, and 0.36 bits in *words6* (see Table 2). In *marbles1*, the mean change was not significantly different from zero, which indicates that participants are probability matching in this condition.

Pairwise comparison of regularization between conditions is also obtained from this re-leveled model (see Table 3). Here, the estimate shows the mean difference in bits for all pairwise comparisons. For example, the mean of *marbles1* was 0.17917 bits lower than that of *marbles1*.

[1] From here onward, whenever we refer to the "entropy of a ratio" we mean the Shannon entropy of the two variants in ratio x:y, where the probability distribution over the variants is $p(V) = \{\frac{x}{10}, \frac{y}{10}\}$.

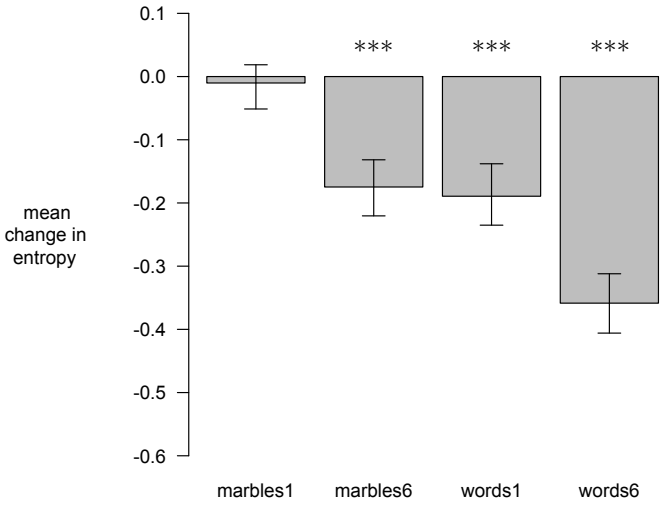

Figure 6: Entropy drops when learners regularize. Each bar shows the average change in Shannon entropy for all pairs of input-output ratios, per condition. Stars indicate significant difference from zero. Error bars indicate the 95% confidence intervals computed with the bootstrap percentile method (Efron, 1979). A significant drop in entropy means that participants regularized in that condition. Non-significant differences from zero are obtained when participants probability match. The lower and upper bounds on mean entropy change for this experiment are -0.67 and $+0.33$ bits.

condition	intercept	S.E.	t statistic	p value
marbles1	-0.01017	0.02904	$t(1152) = -0.350$	$p = 0.73$
marbles6	-0.17474	0.03162	$t(1152) = -5.526$	$p < .001$
words1	-0.18934	0.02904	$t(1152) = -6.519$	$p < .001$
words6	-0.35852	0.03162	$t(1152) = -11.338$	$p < .001$

Table 2: Regularization within each experimental condition.

compare	estimate	S.E.	t statistic	p value
m1, *m6*	-0.16457	0.04293	$t(1152) = -3.833$	$p < .001$
w1	-0.17917	0.04107	$t(1152) = -4.362$	$p < .001$
w6	-0.34835	0.04293	$t(1152) = -8.113$	$p < .001$
m6, *w1*	0.01460	0.04293	$t(1152) = 0.340$	$p = 0.73$
w6	-0.18378	0.04472	$t(1152) = -4.110$	$p < .001$
w1, *w6*	-0.16918	0.04293	$t(1152) = -3.940$	$p < .001$

Table 3: Regularization differences between experimental conditions.

	estimate	S.E.	t statistic
Intercept	0.08567	0.04321	$t(1152) = 1.983$
Load (high)	0.16715	0.04740	$t(1152) = 3.527$
Domain (words)	-0.11373	0.04149	$t(1152) = -2.742$
Input entropy	-0.14094	0.05494	$t(1152) = -2.565$
Load * Input entropy	-0.49568	0.05622	$t(1152) = -8.816$
Domain * Input ent.	-0.09968	0.04596	$t(1152) = -2.169$

Table 4: The of best-fit linear mixed effects model for entropy.

All pairwise comparisons were significantly different, except for that between *words1* and *marbles6*. Overall, participants regularized 26%, 28%, and 53% of the conditional entropy in *marbles6*, *words1*, and *words6*, respectively.

4.3. Domain vs. cognitive load

Effects of the experimental manipulations were assessed by constructing a full linear mixed effects model with three independent variables (i.e. fixed effects) and their interaction: domain, cognitive load, and entropy of the input ratio. The dependent variable was the change in entropy of the input-output ratios. Participant was entered as a random effect (with random intercepts). The significance of each fixed effect was determined by likelihood ratio tests, performed by an ANOVA, on the full model (described above) against a reduced model which omits the effect in question. There was a significant effect of domain ($\chi^2(4) = 46.048, p < .001$), cognitive load ($\chi^2(4) = 105.07, p < .001$), and input ratio ($\chi^2(4) = 520.23, p < .001$). Interactions between fixed effects were also determined by likelihood ratio tests by comparing a reduced model (which omits all interactions) to one which includes the interaction of interest. Two interactions were found to be significant: cognitive load and input ratio ($\chi^2(1) = 74.695, p < .001$) and domain and input ratio ($\chi^2(1) = 4.4462, p = 0.03$). The interaction between domain and cognitive load was not significant ($\chi^2(1) = 0.0059, p = 0.94$).

Therefore, the best-fit model contained an interaction between domain and input ratio, an interaction between cognitive load and input ratio, but only an additive relationship between domain and cognitive load (loglikelihood $= -278.71$). A summary of the best-fit model is given in Table 4. The effect of input ratio on entropy change is due to different amounts of regularization being possible under each input ratio (the maximum drop in entropy achievable under the 5:5 through 0:10 ratios are 1, 0.97, 0.88, 0.72, 0.47, and 0 bits, respectively). As input entropy increases from 0 to 1 bits, output entropy changes with a slope of -0.14094. This means that participants regularize more when the entropy of the input ratio increases from 0 bits (the 10:0 ratio) to 1 bit (the 5:5 ratio). The interactions mean that the slope on input entropy is much steeper (by -0.49568) when cognitive load is high and slightly steeper (by -0.09968) when linguistic stimuli are used. The additive relationship suggests that domain and cognitive load are independent drivers of regularization behavior.

4.4. Frequency-based analysis of regularization

In much of the linguistic regularization literature to date, regularization is measured in terms of stimulus frequency, rather than entropy. In this section, we repeat the analyses from Section 4.2 and 4.3 with a different dependent variable, *change in frequency of the majority variant* (as in, e.g. Hudson Kam & Newport, 2005; Reali & Griffiths, 2009), to illustrate the difference between these two approaches. Figure 7 shows the mean change in frequency of the majority variant (x from input ratio x:y). For example, if a participant produces a 7:3 ratio in response to a 9:1 input ratio, there is a -0.2 change in majority variant frequency for that pair of input-output ratios. In the 5:5 input condition, a random variant was encoded as the

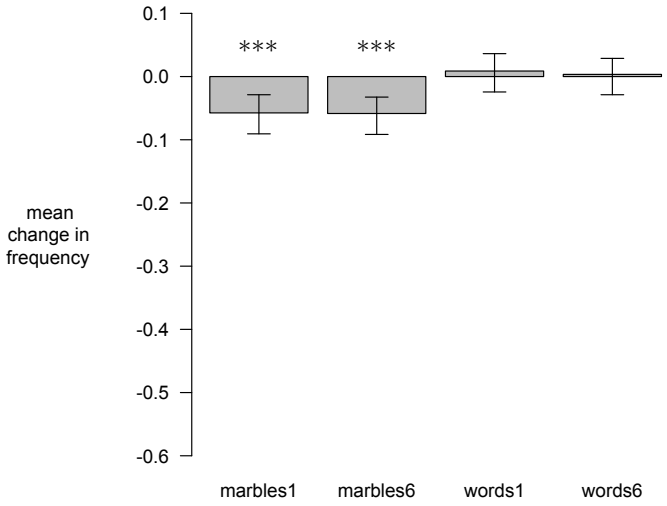

Figure 7: Raw changes in frequency fail to capture regularization behavior. Each bar shows the average difference between the number of times participants observed the majority variant in the training set and the number of times they produced that variant in the testing phase. Error bars indicate the 95% confidence intervals computed with the bootstrap percentile method (Efron, 1979). Values significantly higher than zero indicate a population-level trend in over-producing the majority variant. Values significantly lower than zero indicate a population-level trend in over-producing the minority variant.

condition	intercept	S.E.	t statistic	p value
marbles1	-0.05417	0.01879	$t(1152) = -2.882$	$p = .004$
marbles6	-0.04818	0.01474	$t(1152) = -3.269$	$p = .001$
words1	0.004167	0.01879	$t(1152) = 0.222$	$p = 0.82$
words6	-0.009635	0.01474	$t(1152) = -0.654$	$p = 0.51$

Table 5: Frequency changes within each experimental condition.

	estimate	S.E.	t statistic
Intercept	0.11093	0.04804	$t(1152) = 2.309$
Domain (words)	-0.12580	0.06794	$t(1152) = -1.852$
Input frequency	-0.21521	0.06217	$t(1152) = -3.462$
Domain * Input freq.	0.22925	0.08792	$t(1152) = 2.608$

Table 6: The best-fit linear mixed effects model for frequency.

"majority" variant. Positive changes mean participants over-produced the majority variant and negative changes mean participants over-produced the minority variant.

Applying the analysis in Section 4.2 to the change in majority variant frequency, we find that none of the conditions elicit a significant over-production of the majority variant on average (Table 5), despite the fact that participants in *marbles6*, *words1*, and *words6* are clearly regularizing input ratios (refer back to Figures 5 & 6). However, the frequency-based analysis does reveal something that the entropy-based analysis was unable to capture: a significant over-production of the *minority* variant in the marble-drawing domain, *marbles1* ($t(1152) = -2.882$, $p = .004$) and *marbles6* ($t(1152) = -3.269$, $p = .001$).

For the effects of the experimental manipulations, we apply the analysis in Section 4.3 to the change in majority

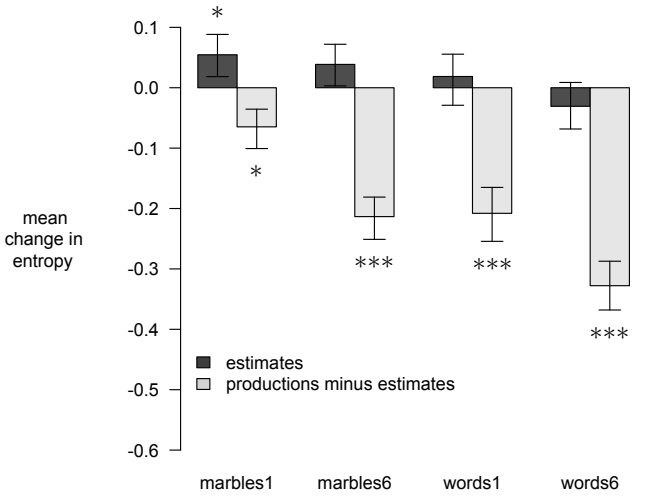

Figure 8: Production bias, not encoding bias, drives regularization. Dark grey: Average difference in regularity between the input ratios participants actually observed and their estimates of the underlying ratio that generated the input ratio. A significant increase in entropy means that participants estimated the underlying ratio to be more variable than the input ratio, and a significant decrease means they estimated it to be more regular. Light grey: Average difference between production ratio regularity and estimated ratio regularity. Error bars indicate the 95% confidence intervals computed with the bootstrap percentile method (Efron, 1979).

variant frequency (and we change the fixed effect *entropy of the input ratio* to *input frequency of the majority variant* in order to match the dependent variable). Table 6 summarizes the best-fit model (loglikelihood −77.0). We find a significant effect of domain ($\chi^2(4) = 16.391, p = 0.003$) and input frequency ($\chi^2(4) = 14.634, p = 0.006$) on change in majority variant frequency, but no significant effect of cognitive load ($\chi^2(4) = 3.0755, p = 0.55$). We also find a significant interaction between domain and input frequency ($\chi^2(4) = 6.7741, p = 0.009$). The effect of input frequency, with a slope of -0.21521, means that participants produce the minority variant more as the input frequency of the majority variant increases (from 5 to 10). However, the interaction between domain and input frequency, with a positive slope of 0.22925, cancels out the effect of input frequency for the linguistic domain. This model can be interpreted as showing an effect of input frequency in the marble-drawing domain only.

In summary, the frequency analysis fails to capture the fact that participants are eliminating variation in the linguistic domain and fails to capture the effect of cognitive load on regularization behavior. The reason mean frequency is not different than zero in the linguistic domain is because participants sometimes regularized with the majority variant and other times regularized with the minority variant, in a way that tends to cause frequency changes to average out to zero. However, as is clear from the raw data, it would be incorrect to conclude that participants are probability matching in the linguistic domain.

condition	intercept	S.E.	t statistic	p value
marbles1	0.05459	0.02388	$t(1152) = 2.286$	$p = 0.02$
marbles6	0.03865	0.02107	$t(1152) = 1.835$	$p = 0.07$
words1	0.01862	0.02388	$t(1152) = 0.780$	$p = 0.44$
words6	−0.03072	0.02107	$t(1152) = -1.458$	$p = 0.15$

Table 7: Mean change in entropy between participants' estimated input ratio and actual input ratio, per experimental condition.

	estimate	S.E.	t statistic
Intercept	0.31193	0.03127	$t(1152) = 9.976$
Load (high)	0.13083	0.03672	$t(1152) = 3.563$
Domain (words)	−0.05145	0.02001	$t(1152) = -2.571$
Input entropy	−0.37040	0.03901	$t(1152) = -9.496$
Load * Input ent.	−0.24258	0.04563	$t(1152) = -5.316$

Table 8: Summary of the best-fit linear mixed effects model.

condition	intercept	S.E.	t statistic	p value
marbles1	−0.06476	0.02535	$t(1152) = -2.554$	$p = 0.01$
marbles6	−0.21339	0.03064	$t(1152) = -6.965$	$p < .001$
words1	−0.20796	0.02535	$t(1152) = -8.202$	$p < .001$
words6	−0.32780	0.03064	$t(1152) = -10.69$	$p < .001$

Table 9: Mean change in entropy between participants' production ratios and estimated input ratios, per experimental condition.

4.5. Regularization during encoding

As discussed in the introduction, regularization behavior is often explained as a result of general cognitive limitations on memory encoding and/or retrieval. The high cognitive load manipulation in this experiment affected both the observation and production phases because both phases consisted of 60 interleaved trials. Therefore, the regularization behavior we observed could be due to encoding multiple frequencies under load (during the observation phase) and/or retrieving frequencies under load (during the production phase). Furthermore, it is possible that linguistic domain may have a specific effect on the encoding of frequency information. To determine whether encoding errors contribute to participants' regularization behavior in this experiment, as described in the method section above we asked participants to estimate (using a slider) the underlying ratio that generated the marble draws or naming events they observed, per container or object (see Section 3.4, last paragraph). If participants' estimates are not significantly different than the ratios they observed, then we can assume frequency encoding was unbiased. This result would point to a production-side driver of regularization.

Figure 8 (dark grey bars) shows the average change in entropy between participants' estimates and the actual input ratios they observed. The same linear mixed effects regression analysis described in Section 4.2 was applied to this data, using the change between input and estimate entropy as the dependent variable (Table 7). Only one condition, *marbles1*, elicited a significant difference between the input ratios and estimates. In this condition, participants estimated the generating ratio to be significantly *more* variable than the ratio they had observed,

indicating a slight encoding bias toward variability. None of the conditions show any bias toward regularity in participants' estimates. Effects of the experimental manipulations were assessed by the same procedure described in Section 4.3, using change between input and estimate entropy as the dependent variable. The best-fit model (loglikelihood = −72.558, Table 8) contained a significant effect of domain ($\chi^2(4) = 11.735, p = 0.02$), cognitive load ($\chi^2(4) = 34.916, p < .001$), and input ratio ($\chi^2(4) = 562.04, p < .001$). One interaction was found to be a significant predictor of participants' estimates: cognitive load and input ratio ($\chi^2(1) = 27.916, p < .001$). Interactions between domain and input ratio ($\chi^2(1) = 0.7554, p = 0.38$) and domain and cognitive load ($\chi^2(1) = 0.6741, p = 0.41$) were not significant. Although the estimate data shows no bias toward regularity, the same factors that affected regularization behavior (cognitive load, domain, and input ratio) also affect participants' estimates. Additionally, we find that the cognitive load manipulation resulted in noisier estimates ($F = 56.487, p < .001$, with Levene's test for homogeneity of variance), whereas the domain manipulation did not ($F = 0.4416, p = 0.51$).

Figure 8 (light grey bars) shows the difference in entropy between the ratio participants produced and their estimate of that ratio, i.e. the extent to which their productions were more regular than their own estimate of their input data. The same linear mixed effects regression analysis described in Section 4.2 was applied to this data, using the difference in entropy between the produced and estimated ratios as the dependent variable (Table 9). In all conditions, production ratios are significantly more regular than the estimates participants made. This means that regularization occurs during the production phase and is likely to be involved in the retrieval and use of frequency information. Interestingly, production-side regularization occurs in all four conditions, even in *marbles1* where participants probability matched their productions to their inputs (effectively "correcting" the variability bias in their estimates). This suggests that regularity is broadly associated with frequency production behavior, even in cases that do not lead to overt regularization behavior.

In summary, raising cognitive load resulted in noisier encoding, however the noise was not biased in the direction of regularity. Estimates in the linguistic domain were not biased toward regularity either. It appears that the bulk of regularization occurs during the production-side of the experiment and is likely to involve processes of frequency retrieval and use.

4.6. Individual differences in frequency learning strategy

The bimodal distributions over output ratios (refer back to Figure 5) suggest individual differences in frequency learning strategies. We break frequency learning behavior into three categories: *regularizing*, *probability matching*, and *variabilizing*. How many participants fall into each category? And in the high load conditions, where participants respond to more than one item, how consistent are their responses across items?

	regularizers	probability matchers	variabilizers
marbles1	10 (5%)	173 (90%)	9 (5%)
words1	50 (26%)	139 (72%)	3 (2%)
marbles6	30 (47%)	15 (23%)	19 (10%)
words6	42 (66%)	14 (22%)	8 (12%)

Table 10: Participants classified by frequency learning strategy. Percentages show how the strategies break down within each condition.

We define *probability matching* as sampling from the input ratio, with replacement. This leads to output ratios that are binomially distributed[2] about the mean (where the mean equals the input ratio). Although the single most likely output ratio a participant could sample is the set of input ratios itself, most probability matchers will sample a ratio that has higher or lower entropy than the input ratio. We will classify participants who produced ratios within the 95% confidence interval of sampling with replacement behavior as probability matchers. We classify participants as *variabilizers* if they produced ratios with significantly higher entropy than likely under probability matching behavior. These could be participants who were attempting to produce a maximally variable set (all 5:5 ratios) or randomly selecting among the two choices on each production trial. Likewise, we classify participants as *regularizers* if they produced ratios with significantly lower entropy than likely under probability matching behavior. It is important to note that a participant with a very weak bias for regularity or variability may consistently produce data that falls within the 95% confidence range of probability matching. However, we take a conservative approach by grouping individuals as regularizers or variabilizers only when probability matching has low probability.

In the low load conditions, where participants only sample one ratio, the 95% confidence intervals on output ratios were determined with the Clopper-Pearson exact method.[3] In the high cognitive load conditions, where participants sample a set of six ratios, we classify the set of ratios according to their conditional entropy $H(V|C)$ (refer back to Section 2). The 95% confidence interval on conditional entropy for probability matching in this experimental setup is 0.43 to 0.75 bits (determined by 10^5 runs of simulated probability matching behavior). Participants who produced data with entropy in the range $0.43 \leq x \leq 0.75$ were classified as probability matchers, those who produced data in the range $0 \leq x < 0.43$ were classified as regularizers, and those who produced data in the range $0.75 < x \leq 1$ were classified as variabilizers.

Table 10 shows the number of participants that fell

[2] Learners can probability match with variance that is significantly lower than binomial variance (Ferdinand, 2015, pp.45-57). Therefore, the definition of probability matching used in this paper is a conservative one.

[3] 95% confidence interval on probability matching per input ratio: 5:5, $0.19 \leq x \leq 0.81$; 6:4, $0.26 \leq x \leq 0.88$; 7:3, $0.35 \leq x \leq 0.93$; 8:2,

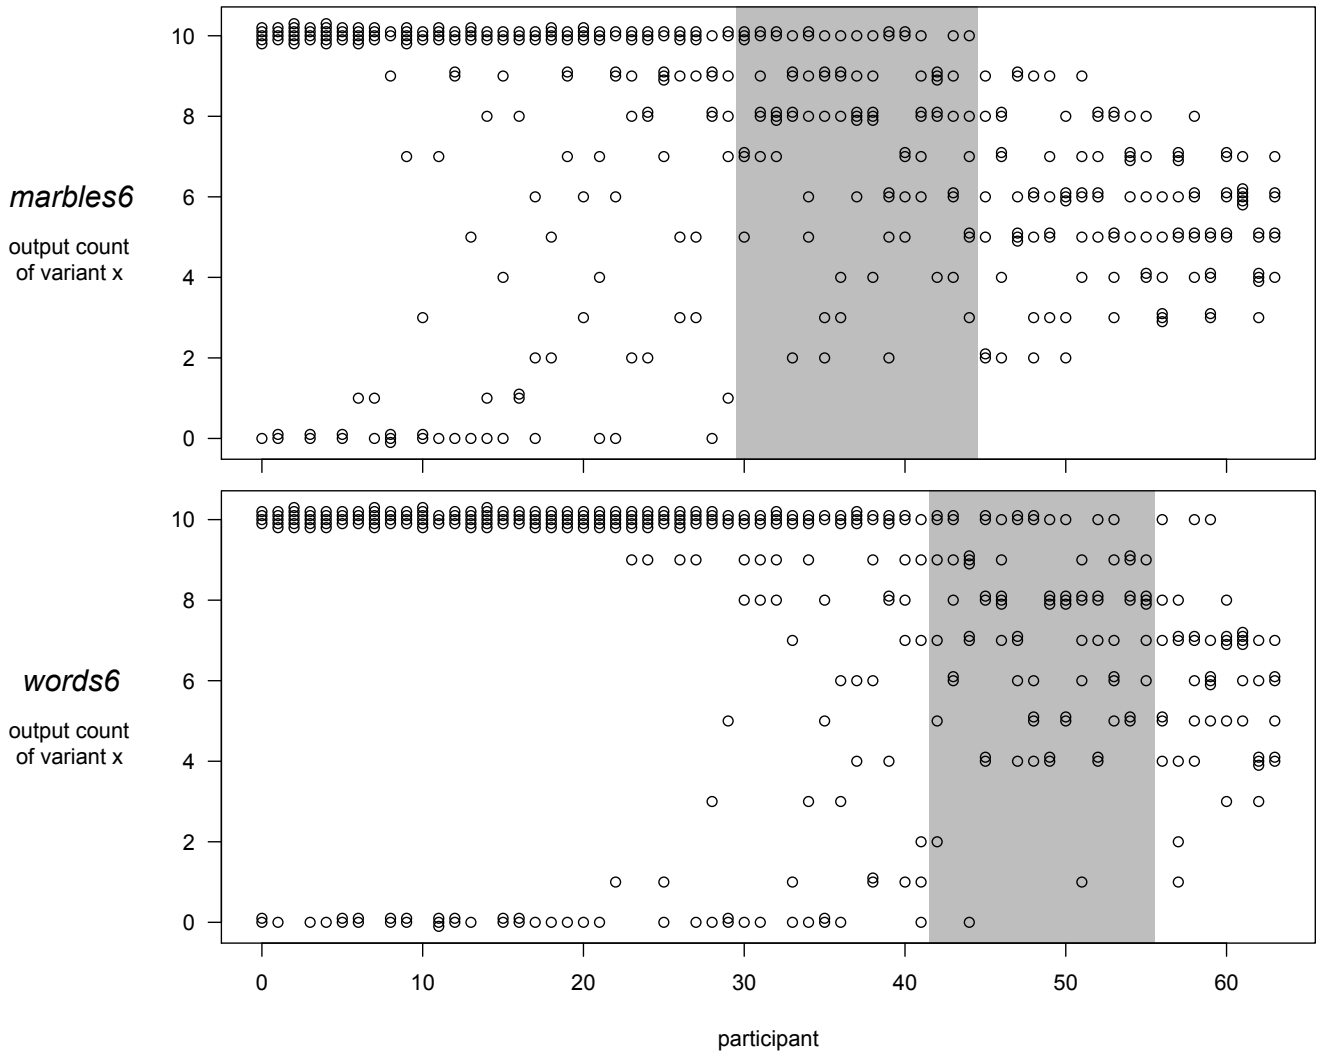

Figure 9: Linguistic and non-linguistic stimuli evoke different frequency learning strategies. Data are from the high cognitive load conditions *marbles6* (top) and *words6* (bottom). The x-axis shows participant number, sorted by their conditional entropy (low to high). The y-axis shows the frequency of the majority variant in the participant's output; each point represents performance on a single container/object, and there are therefore 6 points per participant. The shaded region contains all participants classified as probability matchers. Participants to the left of the shaded region are classified as regularizers and participants to the right are classified as variabilizers.

into each frequency learning category, per condition. All strategies are represented within each experimental condition. There is a significant effect of demand ($\chi^2(2) = 151.63, p < .001$) and domain ($\chi^2(2) = 31.49, p < .001$) on the distribution of frequency learning strategies, meaning that the experimental manipulations elicit different frequency learning strategies from participants. Because fewer data points were collected from participants in the low load condition, probability matching behavior is not easily ruled out, hence the high number of participants classified as probability matchers in *marbles1* and *words1*. It is possible that the difference in dataset size between the low and high conditions is responsible for the significant effect of demand. The effect of domain, however, is reliably due to the experimental manipulation. Therefore, the remainder of this section focuses on the high load data.

Figure 9 shows the set of six output ratios that each participant produced in the high cognitive load conditions. The sets are sorted by their entropy and the shaded box shows the sets that fell into the $0.43 \leq x \leq 0.75$ bit range (classified as probability matchers). Participants to the left of the box are classified as regularizers and participants to the right are classified as variabilizers. More regularizers were found in the linguistic domain, more variabilizers were found in the non-linguistic domain, and probability matchers seem equally likely to be found in either domain. At the extreme left of the x-axis, we see the subset of regularizers, numbering 6 participants in *marbles6* and 22 in *words6*, who produced a maximally regular set (all 10:0 or 0:10, conditional entropy = 0 bits). No participants produced a maximally variable set (all 5:5 ra-

$0.44 \leq x \leq 0.97$; 9:1, $0.55 \leq x \leq 0.99$; 10:0, $0.69 \leq x \leq 1$, where x is the frequency of the majority variant.

tios, conditional entropy = 1 bit). Participants are more likely to maximally regularize in the linguistic condition ($\chi^2(1) = 10.2857, p = .001$).

Points in the 0-4 range on the the y-axis correspond to output ratios that contained a large number of minority variant productions (i.e. the majority variant had frequency of between 0 and 4). In Figure 9 we see that there are *no* participants who exclusively regularized with the minority variant. A few participants regularized with the majority variant exclusively, however most participants regularized with 1-2 minority variants and 4-5 majority variants. In other words, there appear to be no stable individual differences in tendency to regularize with the minority variant.

4.7. Primacy and recency effects on regularization

Studies on regularization often find that participants regularize by over-producing or over-predicting the majority variant, and this serves as the standard definition of regularization (e.g. Hudson Kam & Newport, 2005, 2009). However, several studies, including this one, report some participants who regularize with the minority variant (e.g. Reali & Griffiths, 2009; Smith & Wonnacott, 2010; Culbertson et al., 2012; Perfors, 2012, 2016). What causes some participants to regularize with the majority variant, and others to regularize with the minority variant? In the previous section, we saw minority regularization is not due to individual differences in frequency learning behavior. If minority regularization is not a feature of individuals, it may be a feature of the training data they received.

One possible data-driven explanation for minority regularization lies in the effects of a stimulus's primacy and recency on participant behavior. In the observation phase, participants were presented with a randomly-ordered sequence of variants, such that the probability of any particular variant occurring at the beginning or end of the input sequence is proportional to its frequency in the sequence. Therefore, some participants would have received minority variants toward the beginning and/or end of the sequence, whereas others would have not. Many experiments on the serial recall of lexical items show that participants are better at recalling the first and last few items in a list of words (e.g. Deese & Kaufman, 1957; Murdock, 1962). This effect also extends to the learning of mappings between words and referents: Poepsel & Weiss (2014) found that when participants in a cross-situational learning task were confronted with several possible synonyms for an object, their confidence in a correct mapping was positively correlated with the primacy of that mapping in the observation phase. Therefore, we investigated the effect of the minority variant's position in the input sequence on participants' tendency to regularize with the minority variant.

Unlike most research on primacy and recency (which present participants with a long list of unique stimuli), our input sequences only consist of two variants, presented several times each. Therefore, we can quantify the strength of minority primacy as the imbalance of the variants across

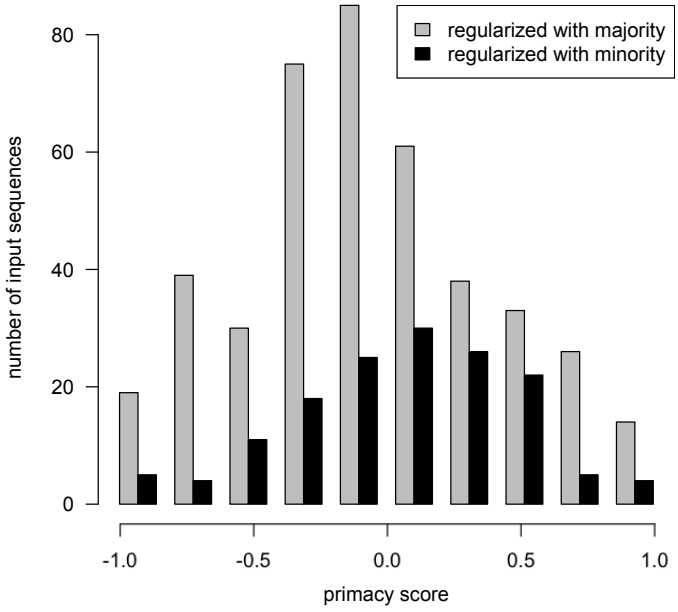

Figure 10: Participants are more likely to regularize with the minority variant when they observe it toward the beginning of the input sequence. The x-axis is the primacy of the minority variant in the input sequence, ranging from -1 (maximal recency) to 1 (maximal primacy). Bars show the number of input sequences that were regularized by over-producing the minority variant (black) and by over-producing majority variant (grey).

production sequences	*marbles1*	*words1*	*marbles6*	*words6*
total	192	192	384	384
regularized	43	85	201	241
regularized w/minority	16 (37%)	18 (21%)	53 (26%)	63 (26%)

Table 11: Number of regularized production sequences per condition. Parentheses show the number of minority-regularized sequences as a percentage of all regularized sequences.

the input sequence. To do this, we will use the notion of net torque. In this analogy, we consider the input sequence to be a weightless lever of length 10 (the number of observation trials), we consider each minority variant to be a weight of one unit which is placed on the lever according to its observation trial number, and we assume the lever is balanced on a fulcrum at its center. The sum of the distance of the weights located right of center minus the sum of the distance of the weights left of center is the net torque. We will use the following standardization of net torque[4], and refer to it as the primacy score:

$$primacy\ score = -\left(\frac{\sum_{d=1}^{N} w_d(d - \frac{N+1}{2})}{m(N-m)/2}\right) \tag{4}$$

where w is the sequence of weights and d is the distance of that weight from the start of the sequence. In the 5:5

[4]Thanks to Andrew Berdahl for providing this solution. The score was standardized in order to de-correlate net torque with input ratio.

input sequences, a random variant is coded as the "minority" variant. N is the length of w and m is the total number of minority variants in the sequence. Positive values mean that the minority variants occur more toward the beginning of the sequence and negative values mean they occur more toward the end of the sequence. The maximum primacy score is 1 and the minimum is -1. The average primacy score is 0 and is obtained when the sequence is balanced (i.e. minority variants are equally distributed early and late in the input sequence). For example (where 1 indicates an occurrence of the minority variant in the input sequence), the primacy score of sequence 1110000000 is 1, 0000000001 is -1, 0101001000 is 0.33, 1000000001 is 0, and 0000110000 is 0.

Table 11 shows a breakdown of the number of regularized production sequences per experimental condition (i.e. all output sequences that had lower entropy than their corresponding input sequence). Analyses were restricted to the 570 input sequences that participants regularized. Figure 10 plots the primacy scores of the 420 sequences that were regularized with the majority variant (grey) and the 150 sequences that were regularized with the minority variant (black). We constructed a logit mixed effects model of regularization type (majority or minority regularization) as a function of primacy score. Participant was entered as a random effect (with random intercepts). A likelihood ratio test was performed by an ANOVA on this model and a reduced model which omits primacy score as a predictor. We found a significant effect of primacy score on regularization type ($\chi^2(1) = 6.4082, p = 0.01$). On average, primacy score is 0.11 points higher ($\pm$ 0.04 standard errors) in sequences that were regularized with the minority. This means that participants are more likely to regularize with the minority variant when they saw it toward the beginning of their input sequence (i.e. when minority variant primacy is high). However, minority regularization is not entirely explained by minority primacy. As can be seen in Figure 10, minority regularization was obtained across all primacy scores and even when the minority was maximally recent (left-most black bar).

5. Predicting the evolution of regularity

In the previous sections, we showed that learners regularize frequencies due to domain-general and domain-specific constraints. This was accomplished by analyzing *one cycle of learning*, which spans the perception, processing, and production of a set of linguistic variants. Although this informs us about the relevant constraints that underwrite regularity in language, and even *how much* regularity each constraint imposes, it does not necessarily tell us how much regularity we will expect to see in a language over time. This is because languages are transmitted between generations of learners and are therefore subject to multiple learning cycles, where each individual has an opportunity to impose some amount of regularity on the language. We have seen that several factors (domain, cognitive load, input ratio, individual differences, and primacy effects) create a complex landscape for the evolution of regularity. How does regularity accumulate in a language over time? Is it a simple, compounding effect and should we expect all linguistic variation to be eliminated eventually? Or will languages converge on a certain level of regularity?

In this section, we provide some perspective on these questions by analyzing the dynamics of change in our artificial language data. One way to explore the dynamics of language evolution is through iterated learning (Kirby et al., 2014) in which the output of one learner serves as the input to another (e.g. Kirby, 2001; Brighton, 2002; Smith et al., 2003; Kirby et al., 2008; Reali & Griffiths, 2009; Smith & Wonnacott, 2010). Several cycles of iterated learning result in a walk over the complex landscape of constraints that shape language, and several walks can be used to estimate this landscape and likely evolutionary trajectories. Griffiths & Kalish (2007) have shown that iterated learning is equivalent to a Markov process, which is a discrete-time random process over a sequence of values of a random variable, $v_{t=1}, v_{t=2}, ..., v_{t=n}$, such that the random variable is determined only by its most recent value (Papoulis, 1984, p.535):

$$P(v_t|v_{t=1}, v_{t=2}, ..., v_{t-1}) = P(v_t|v_{t-1}) \quad (5)$$

This describes a memoryless, time-invariant, process in which the past values $(v_{t-2}, v_{t-3}, ..., v_{t-n})$ of the variable have no direct influence on the current value (v_t). This is the case for iterated learning chains when learners only observe the behaviors of the previous generation. All of the possible values of the random variable constitute the state space of this system. A Markov process is fully specified by the probabilities with which each state will lead to every other state and these probabilities between states can be represented as a transition matrix, $\mathbf{Q}$ (Norris, 2008, p.3). The probabilities in $\mathbf{Q}$ are the aforementioned landscape over which a language evolves.

In the experimental data, each state s corresponds to one of the eleven possible ratios: $s_0, s_1, ..., s_{10} = \{$0:10, 1:9, 2:8, 3:7, 4:6, 5:5, 6:4, 7:3, 8:2, 9:1, and 10:0$\}$, where s_{t-1} is the input ratio and s_t is the output ratio. Our experiment was designed so that $\mathbf{Q}$ could be estimated for each of the four experimental conditions, by collecting data from participants in each of the eleven possible states. Figure 11 (top row) shows the estimated transition matrix from each experimental condition. Each estimation consists of the raw data in that condition, smoothed with a small value $\epsilon = \frac{1}{length(row)^2}$. Each cell in the matrix, $\mathbf{Q}_{ij}$, gives the transition probability from state $s_{i=t-1}$ to state $s_{j=t}$. The shading of the cells denote the transition probabilities between states. Each row in the matrix corresponds to the distribution of output ratios produced in response to one input ratio (rows are the same distributions in Figure 5, only smoothed). For example, row 5 in the *marbles1* transition matrix corresponds to the upper left panel of Figure

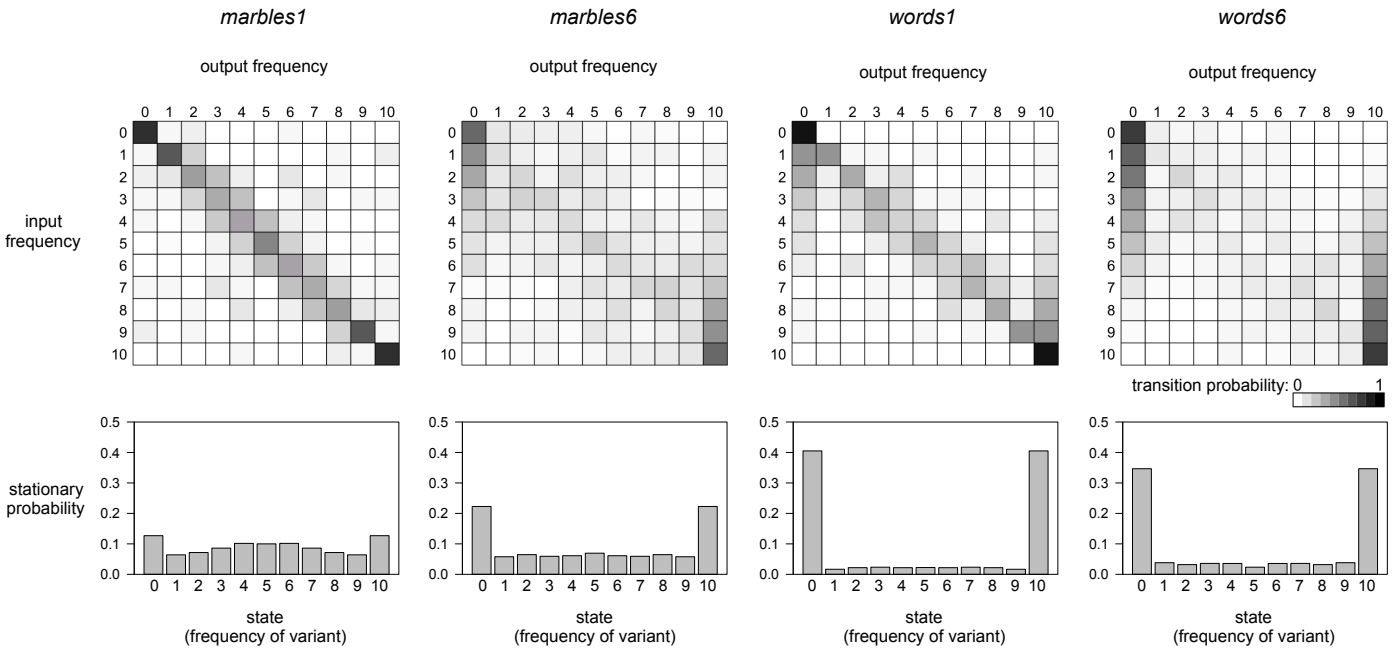

Figure 11: The data from the experiment is used to predict the cultural evolution of regularization. Top: Estimated transition matrices for each experimental condition contain the probabilities that a learner produces any given output ratio from any given input ratio (presented in terms of the frequency of variant x in each input ratio x:y). Bottom: The stationary distribution shows the percentage of learners who will produce each output ratio, after the ratios have evolved for an arbitrarily large number of generations. Each stationary distribution is the solution to the matrix above it.

5, and the probability of transitioning from $s_{t-1=5}$ to $s_{t=6}$ is equivalent to the (smoothed) proportion of participants that produced a 6:4 ratio when trained on a 5:5 ratio. Likewise, rows 4 and 6 correspond to the 6:4 panel in Figure 5, but this distribution is flipped in row 4 to display the results in terms of the minority variant.

The transition matrices can be used to estimate the regularity of the data after an arbitrarily large number of learning cycles. No matter what start state is used to initialize an iterated learning chain, an arbitrarily large number of iterations will converge to a stationary distribution, $\vec{s}$. The stationary distribution is defined as $\vec{s}\mathbf{Q} = \vec{s}$, meaning that once the data take the form of the stationary distribution and serve as the input to $\mathbf{Q}$, the output will be the same distribution and the subsequent generations of data will not change anymore. The stationary distribution is a probability distribution over all states in the system, where each probability corresponds to the proportion of time the system will spend in each state, and can be solved for any matrix by decomposing the matrix into its eigenvalues and eigenvectors: $\vec{s}$ is proportional to the first eigenvector. Figure 11 (bottom row) shows the stationary distribution for each transition matrix. From these distributions, for example, we see that an arbitrarily long iterated learning chain will produce maximally regular (0:10 and 10:0) ratios approximately 25% of the time if participants are learning about two marbles and one container (*marbles1*) and approximately 80% of the time when participants are learning about two words for one object (*words1*).

We calculate the level of regularity in the stationary distribution by multiplying the Shannon entropy of the ratio (defined by each state, $\vec{s}_i$) by the probability of observing that state, $p(\vec{s}_i)$. The results are 0.61 bits of conditional entropy $H(V|C)$ in *marbles1*, 0.43 bits in *marbles6*, 0.16 bits in *words1*, and 0.24 bits in *words6*. We compare these to the results of the experiment (the average conditional entropy achieved after one learning cycle), which was 0.66 bits in *marbles1*, 0.50 bits in *marbles6*, 0.48 bits in *words1*, and 0.32 bits in *words6*.

Figure 12 plots the results above in terms of entropy change: as the difference between the mean input entropy used in the experiment and the mean entropy after one learning cycle (in dark grey) and after convergence to the stationary distribution (in light grey). First, we find that variation will never be completely lost: the data in each experimental condition converges on a certain level of regularity rather than a maximally regular system. A maximally regular system would have a conditional entropy of 0 bits and an entropy change of -0.67 bits in this experimental setup, but none of the conditions reach this lower bound.

Second, we find that regularity appears to increase over the course of cultural transmission and in the case of *words1*, the increase is significant (inferred from non-overlapping 95% confidence intervals). This result, where cultural transmission increases the effect of the bias on the data, has been demonstrated and discussed by Kalish et al.

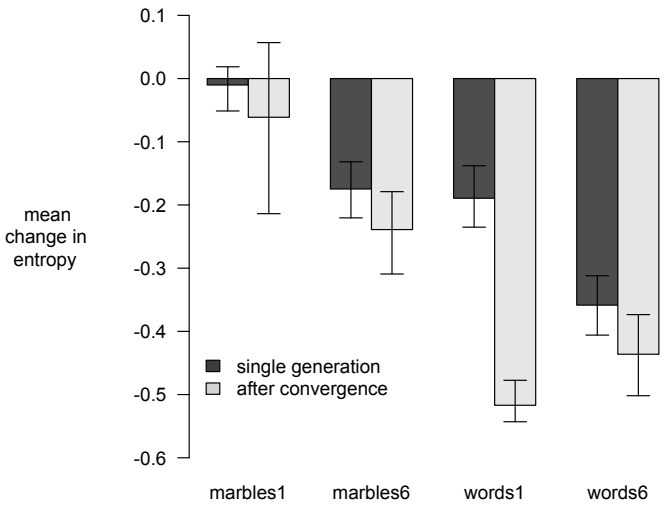

Figure 12: Same learning biases lead to different degrees of regularization after many generations of cultural transmission. Dark grey: Average change in entropy after one learning cycle (same data in Figure 6, reprinted here for comparison). Light grey: Average change in entropy of variants after convergence to the stationary distribution (i.e. after an infinite number of learning cycles). Error bars indicate 95% confidence intervals, computed by the bootstrap percentile method (Efron, 1979) on 10,000 resamples of the transition matrix, where each matrix was solved for its stationary distribution and mean change in entropy.

(2007); Griffiths et al. (2008); Reali & Griffiths (2009) in terms of biases being revealed by cultural transmission, and by Kirby et al. (2007); Griffiths & Kalish (2007); Smith & Wonnacott (2010); Thompson (2015) in terms of cultural transmission amplifying weak biases. As for regularization biases specifically, Reali & Griffiths (2009) showed that the regularity of a culturally transmitted artificial language increases in just a few generations.

Third, we find that different regularization biases are not amplified at similar rates. Although *marbles6* and *words1* yield similar amounts of regularity after one generation, these amounts differ markedly after many generations. This difference is due to the different distribution of probabilities within the transition matrix, which can attract iterated learning chains to different regions of the state space. One reason why the *words1* data regularizes more than the other data sets, is that it has a markedly lower probability of transitioning out of the 10:0 and 0:10 states, trapping generations of learners in this highly regular region for longer amounts of time.

Our third finding has important implications for the relationship between learning biases and structure in language: it means that culturally transmitted systems, such as language, do not necessarily simply mirror the biases of its learners (see Kirby, 1999; Kirby et al., 2004; Smith et al., 2017). Previously, we showed that cognitive load and linguistic domain are independent sources of regularization in individual learners. Looking at the data from individual learners, we may even infer that cognitive load and linguistic domain inject similar amounts of regularity into language. However, this does not mean that a data set which is culturally transmitted under conditions of only high cognitive load (as in *marbles6*) or only linguistic framing (as in *words1*) will ultimately acquire the same amount of regularity. The fact that *words1* has higher stationary regularity than *marbles6* means, at least in terms of the present data, that the amount of regularity we ultimately expect to find in a language is not simply predicted from the learning biases. Instead, the process of cultural transmission is an indispensable piece of the puzzle in explaining how learning biases shape languages.

6. Discussion

Regularity in language is rooted in the cognitive apparatus of its learners. In this paper, we have shown that linguistic regularization behavior results from at least two, independent sources in cognition. The first is domain-general and involves constraints on frequency learning when cognitive load is high. The second is domain-specific and is triggered when the frequency learning task is framed with linguistic stimuli.

Cognitive load was manipulated by varying the number of stimuli in a frequency learning task. When participants observed and produced for more stimuli, they regularized stimuli frequencies *more* on average than when they were observing and producing for fewer stimuli. This result held when stimuli were non-linguistic (marbles and containers) and when stimuli were linguistic (words and objects) and has previously been observed in separate non-linguistic and linguistic experiments (Gardner, 1957; Hudson Kam & Newport, 2009). We have shown, within the same experimental setting and for identical distributions of variation, that increasing cognitive load causes participants regularize both non-linguistic and linguistic stimuli. Furthermore, we have shown that participants regularize a similar amount of variation in both cases, eliminating 24.6% of the variation in marbles conditioned on containers and 25.5% of the variation in words conditioned on objects. This similarity suggests that learners have general limits on the amount of variation they can process and reproduce, which are independent of the learning domain. It is quite possible that cognitive load makes a fixed contribution to regularization behavior, however it remains to be seen whether this result holds over a variety of other distributions with different entropy levels.

Domain was manipulated by varying the type of stimuli used in the frequency learning task. When participants observed and produced mappings between words and objects, they regularized more than participants who observed and produced mappings between marbles and containers. Participants appear to have a higher baseline regularization behavior when learning about linguistic stimuli: an additional 27% of variation was regularized due to linguistic domain in each cognitive load condition (26.7% in the low condition, 27.4% in the high condition). The use of linguistic stimuli may trigger any number of domain-specific

learning mechanisms or production strategies. One possibility is that the stimuli manipulation changed participants' pragmatic inferences about the frequency learning task. In an artificial language learning task, Perfors (2016) showed that participants regularize more when they believe that the variation in labels for objects can be the result of typos, suggesting that participants are more likely to maintain variation when they think it is meaningful. It is possible that participants make different assumptions about the importance of variation in marbles versus words when they are required to demonstrate what they have learned to an experimenter. However, it is not clear what these assumptions may be. Another possibility is that the use of linguistic stimuli encourages participants to consider the communicative consequences of variation. Participants in artificial language learning tasks regularize more when they are allowed to communicate with one another (Fehér et al., 2016; Smith et al., 2014) and even when they erroneously believe they are communicating with another participant (Fehér et al., 2016). This suggests that participants strategically regularize variation in situations that are potentially communicative and may be the reason that regularization is observed in a wide range of language learning tasks.

We also investigated the role of encoding errors on regularization behavior. After the production phase, participants estimated the ratio of the variants associated with each container or object they had observed. We found that domain and cognitive load affect estimates: more variation is encoded when cognitive load is low and stimuli are non-linguistic. However, the estimates themselves were not significantly more regular than the input ratios that participants observed. This suggests that participants had access to somewhat accurately encoded frequency information when making their estimates. However, it is possible that biased encoding could result from more complex mappings than those used in this experiment. Vouloumanos (2008) found that learners are able to encode and retrieve fine-grained differences in the statistics of low-frequency mappings between words and objects (which we calculate had a joint entropy of 4.41 bits), but failed to encode and retrieve fine-grained differences for a more complex stimuli set (with a joint entropy of 5.15 bits). The joint entropy of our high cognitive load mappings, at 3.26 bits, is within Vouloumanos (2008)'s demonstrated threshold for accurate frequency representation. Our finding relates to the Less-is-More hypothesis (Newport, 1990), which describes regularization as a result of memory limitations on frequency learning. Under this hypothesis, participants may fail to encode, store, or retrieve lower-frequency linguistic forms, and therefore fail to produce and perpetuate these forms in their language (effectively reducing the variation in the language). Although our cognitive load manipulation resulted in noisier encoding when load was high (showing that the high load frequency learning task was indeed more difficult for participants), that noise was not biased toward more regular ratios. Because participants regularized their productions without showing a corresponding bias in estimates, this implies that the bulk of their regularization occurred during the production phase of the task. This production-side interpretation is in line with the results of Hudson Kam & Chang (2009), who showed that participants regularize more when stimuli retrieval is made harder, and Perfors (2012), who found participants do not regularize when encoding is made harder. Taken together, these results suggest that the Less-is-More hypothesis applies more to retrieval and less to encoding. Furthermore, given our observation that many participants regularize with the minority variant, retrieval errors may not be exclusively based on the failure to access lower-frequency forms.

In a further exploration of minority regularization behavior, we found that it was not the result of individual differences in regularization behavior. In the high cognitive load condition, where participants responded to 6 items, no participant regularized with the minority variant on all six items. Although participants differed in frequency learning strategies (we found regularizers, probability matchers, and variabilizers in all four conditions), most participants tended to regularize with one or two minority variants. If minority regularization is not the product of differences in individual frequency learning strategies, it could be the product of differences in the randomized stimuli that each participant saw. Therefore, we investigated the role of minority variant primacy in the observation phase and found that participants are slightly more likely to regularize with the minority variant when it occurs toward the beginning of the observation sequence. This finding is in line with the results of Poepsel & Weiss (2014), who showed that participants in a cross-situational learning task had higher confidence in the correctness of a mapping between words and referents when those items co-occurred early in the observation phase.

Another issue surrounding minority regularizers is that they can confound regularization analyses which are based on the majority variant's change in frequency. Alternative analyses that overcome this issue are Perfors (2012, 2016)'s *regularization index* and entropy-based analyses, as we use here. Regularization should not be defined exclusively as "overproduction of the highest-frequency or dominant form". Regularization occurs whenever learners increase the predictability of a linguistic system and therefore equates to a system's decrease in entropy. Overproduction of dominant forms certainly can cause a language's entropy to drop, however regularity also can increase when minority forms are overproduced or when forms are maintained but conditioned on other linguistic contexts or meanings. We found that entropy and frequency analyses are sensitive to different aspects of the data. Entropy is better for quantifying regularization and positively identifying it, whereas frequency is better for detecting a population-level trend in over- or under-producing a particular variant. For example, the frequency method did not capture the effect of cognitive load on frequency learning behavior,

but it did capture an interesting domain difference in this experiment: marble drawers overproduced the minority variant on average, whereas word learners did not. These two methods also show differences in the classification of probability matching behavior: the entropy method identified *marbles1* as consistent with probability matching behavior and the frequency method did not (because there is a significant bias toward the minority variant). This raises important questions about the nature of probability matching: should it be defined as reproducing the same amount of variation (as the entropy measure captures) or reproducing the same amount of variation *along with* the correct mapping of variation to stimuli (as the frequency measure captures)?

Overall, this paper explored how various cognitive constraints on frequency learning give rise to regularization behavior. But what can detailed knowledge of these constraints tell us about the regularity of languages? One possibility is that the relationship between constraints on learning and structure of languages is straightforward, such that constraints or biases in learning can be directly read off the typology of languages in the world (e.g. Baker, 2001). Griffiths & Kalish (2007) also find a simple relationship in a Bayesian model of cultural transmission for learners who sample hypotheses from the posterior distribution: the probability that a learner ends up speaking any given language is equal to the probability of the language in the prior (i.e. the learner's learning bias). Under many other conditions, however, cultural transmission distorts the effects of learners' biases on the data they transmit, making it impossible to simply read learning biases off of language universals (Kirby, 1999; Kirby et al., 2004). Sometimes this distortion is one of amplification, where the effects of the bias accumulate, such that weak biases have strong effects on the structure of culturally transmitted data (e.g. Kirby et al., 2007; Smith & Wonnacott, 2010; Thompson, 2015). However, the opposite can also occur: biases can also have weaker effects or no effects at all (Smith et al., 2017). This suggests that cultural transmission increases the complexity of the relationship between individual learning biases and the structure of language. In this paper, we investigated the relationship between regularization biases and regularity in culturally transmitted datasets by plugging the data obtained from our population of participants into a model of cultural transmission. We found that the relative contribution of two sources of regularization bias, increased cognitive load and linguistic domain, changes in a complex way over time. Regularity in the culturally transmitted data increased (i.e. was amplified) over time, however the rates of amplification were different for load and domain: linguistic domain contributed more to regularity over time, but an interaction lessened this effect when cognitive load was high. From this experiment, we found that several factors certainly affect regularization behavior: domain, cognitive load, input ratio, individual differences, and stimuli primacy. And the modeling work shows, within the same experimental data set, how the apparent relative contributions of biases to behavior can change via cultural transmission. Understanding the effects of cognitive biases on human behavior is a challenging task, but an even greater challenge is posed in understanding the effects of biases on behavior that evolves culturally.

7. Conclusion

When learners observe and reproduce probabilistic variation, we find they regularize (reduce variation) when cognitive load is high and when stimuli are linguistic. We conclude that regularity in natural languages is a co-product of domain-general and domain-specific biases on frequency learning and production. Furthermore, we find that load and domain affect how participants encode frequency information. However, encoded frequencies are not more regular than the data participants observed: the bulk of regularization occurs when participants produce data. Finally, we show that the relative contributions of load and domain to the regularity in a language can change when data are transmitted culturally. In order to understand how various regularity biases create regularity in language, experiments that quantify learning biases need to be coupled with cultural transmission studies.

Acknowledgements

This research was supported by the University of Edinburgh's College Studentship, the SORSAS award, and the Engineering and Physical Sciences Research Council. The writeup was supported by the Omidyar Fellowship. We thank Bill Thompson, Simon DeDeo, Tom Griffiths, Luke Maurits, and the members of the Centre for Language Evolution for their feedback on this project. We thank Florencia Reali and Daniel Richardson for providing experimental stimuli.

References

References

Baker, M. C. (2001). *The atoms of language: The mind's hidden rules of grammar*. New York: Basic books.

Bates, D., Maechler, M., Bolker, B., & Walker, S. (2013). *lme4: Linear mixed-effects models using Eigen and S4*. URL: `http://cran.r-project.org/package=lme4`.

Becker, A., & Veenstra, T. (2003). The survival of inflectional morphology in French-related creoles. *Studies in Second Language Acquisition*, *25*, 283–306.

Berko, J. (1958). *The child's learning of English morphology*. Ph.D. thesis Radcliffe College.

Bickerton, D. (1981). *Roots of Language*. Ann Arbor, MI: Karoma Publishers.

Brighton, H. (2002). Compositional syntax from cultural transmission. *Artificial life*, *8*, 25–54.

Chomsky, N. (1957). *Syntactic Structures*. Mouton Publishers.

Cover, T. M., & Thomas, J. A. (1991). *Elements of Information Theory*.

Culbertson, J., & Kirby, S. (2015). Simplicity and specificity in language: Domain-general biases have domain-specific effects. *Frontiers in psychology*, *6*.

Culbertson, J., & Newport, E. L. (2015). Harmonic biases in child learners: In support of language universals. *Cognition*, *139*, 71–82.

Culbertson, J., Smolensky, P., & Legendre, G. (2012). Learning biases predict a word order universal. *Cognition*, *122*, 306–29.

Cuskley, C., Colaiori, F., Castellano, C., Loreto, V., Pugliese, M., & Tria, F. (2015). The adoption of linguistic rules in native and non-native speakers: Evidence from a wug task. *Journal of Memory and Language*, *84*, 205–223.

Deese, J., & Kaufman, R. (1957). Serial effects in recall of unorganized and sequentially organized verbal material. *Journal of experimental psychology*, *54*, 180–187.

DeGraff, M. (1999). Creolization, language change and language acquisition: An epilogue. In M. DeGraff (Ed.), *Language creation and language change: Creolization, diachrony, and development* (pp. 473–543). Cambridge, MA: MIT Press.

Dougherty, M. R. P., & Hunter, J. (2003). Probability judgment and subadditivity: The role of working memory capacity and constraining retrieval. *Memory & Cognition*, *31*, 968–982.

Efron, B. (1979). Bootstrap methods: another look at the jackknife. *The annals of Statistics*, (pp. 1–26).

Fedzechkina, M. (2014). *Communicative efficiency, language learning, and language universals*. Ph.D. thesis University of Rochester.

Fehér, O. (2016). Atypical birdsong and artificial languages provide insights into how communication systems are shaped by learning, use, and transmission. *Psychonomic Bulletin & Review*, (pp. 1–9).

Fehér, O., Wonnacott, E., & Smith, K. (2016). Structural priming in artificial languages and the regularisation of unpredictable variation. *Journal of Memory and Language*, *91*, 158–180.

Ferdinand, V. A. (2015). *Inductive evolution: cognition, culture, and regularity in language*. Ph.D. thesis The University of Edinburgh.

Gardner, R. A. (1957). Probability-learning with two and three choices. *The American Journal of Psychology*, *70*, 174–185.

Givón, T. (1985). Function, structure, and language acquisition. In D. Slobin (Ed.), *The crosslinguistic study of language acquisition.* (pp. 1005–1028). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Griffiths, T., & Kalish, M. (2007). Language evolution by iterated learning with Bayesian agents. *Cognitive Science*, *31*, 441–480.

Griffiths, T. L., Christian, B. R., & Kalish, M. L. (2008). Using category structures to test iterated learning as a method for identifying inductive biases. *Cognitive Science*, *32*, 68–107.

Hauser, M. D., Chomsky, N., & Fitch, W. T. (2002). The faculty of language: what is it, who has it, and how did it evolve? *science*, *298*, 1569–1579.

Hudson Kam, C. L., & Chang, A. (2009). Investigating the cause of language regularization in adults: memory constraints or learning effects? *Journal of experimental psychology. Learning, memory, and cognition*, *35*, 815–21.

Hudson Kam, C. L., & Newport, E. L. (2005). Regularizing unpredictable variation: the roles of adult and child learners in language formation and change. *Language learning and development*, *1*, 151–195.

Hudson Kam, C. L., & Newport, E. L. (2009). Getting it right by getting it wrong: when learners change languages. *Cognitive psychology*, *59*, 30–66.

Kalish, M. L., Griffiths, T. L., & Lewandowsky, S. (2007). Iterated learning: Intergenerational knowledge transmission reveals inductive biases. *Psychonomic Bulletin & Review*, *14*, 288–294.

Kareev, Y., Lieberman, I., & Lev, M. (1997). Through a narrow window: Sample size and the perception of correlation. *Journal of Experimental Psychology*, *126*, 278–287.

Kirby, S. (1999). *Function, selection and innateness: the emergence of language universals*. Ph.D. thesis University of Edinburgh.

Kirby, S. (2001). Spontaneous evolution of linguistic structure-an iterated learning model of the emergence of regularity and irregularity. *Evolutionary Computation, IEEE Transactions on*, *5*, 102–110.

Kirby, S., Cornish, H., & Smith, K. (2008). Cumulative cultural evolution in the laboratory: An experimental approach to the origins of structure in human language. *PNAS*, *105*, 10681–10686.

Kirby, S., Dowman, M., & Griffiths, T. (2007). Innateness and culture in the evolution of language. *Proceedings of the National Academy of Sciences*, *104*, 5241–5245.

Kirby, S., Griffiths, T., & Smith, K. (2014). Iterated learning and the evolution of language. *Current opinion in neurobiology*, *28*, 108–114.

Kirby, S., Smith, K., & Brighton, H. (2004). From ug to universals: Linguistic adaptation through iterated learning. *Studies in Language*, *28*, 587–607.

Lieberman, E., Michel, J.-B., Jackson, J., Tang, T., & Nowak, M. A. (2007). Quantifying the evolutionary dynamics of language. *Nature*, *449*, 713–716.

Lumsden, J. S. (1999). Language creation and language change. In M. DeGraaf (Ed.), *Language acquisition and creolization*. Cambridge, MA: MIT Press.

Marcus, G. F., Pinker, S., Ullman, M., Hollander, M., Rosen, T. J., & Xu, F. (1992). Overregularization in Language Acquisition. *Monographs of the Society for Research in Child Development*, *57*, i–178.

Meyerhoff, M. (2000). The emergence of creole subject–verb agreement and the licensing of null subjects. *Language Variation and Change*, *12*, 203–230.

Murdock, B. B. (1962). The serial position effect of free recall. *Journal of experimental psychology*, *64*, 482.

Newport, E. L. (1990). Maturational constraints on language learning. *Cognitive science*, *14*, 11–28.

Norris, J. R. (2008). *Markov chains*. Cambridge University Press.

Papoulis, A. (1984). Brownian movement and markoff processes. *Ch*, *15*, 515–553.

Pearl, L., & Lidz, J. (2009). When domain-general learning fails and when it succeeds: Identifying the contribution of domain specificity. *Language Learning and Development*, *5*, 235–265.

Perfors, A. (2012). When do memory limitations lead to regularization? An experimental and computational investigation. *Journal of Memory and Language*, (pp. 1–60).

Perfors, A. (2016). Adult regularization of inconsistent input depends on pragmatic factors. *Language Learning and Development*, *12*, 138–155.

Poepsel, T. J., & Weiss, D. J. (2014). Context influences conscious appraisal of cross situational statistical learning. *Frontiers in psychology*, *5*.

R Core Team (2013). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing Vienna, Austria. URL: `http://www.r-project.org/`.

Reali, F., & Griffiths, T. L. (2009). The evolution of frequency distributions: relating regularization to inductive biases through iterated learning. *Cognition*, *111*, 317–28.

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, *274*, 1926–1928.

Saffran, J. R., & Thiessen, E. D. (2007). Domain-general learning capacities. *Blackwell handbook of language development*, (pp. 68–86).

Sankoff, G. (1979). The Genesis of a Language. In K. C. Hill (Ed.), *The Genesis of Language* (pp. 23–47). Ann Arbor, MI: Karoma Publishers.

Schilling-Estes, N., & Wolfram, W. (1994). Convergent explanation and alternative regularization patterns: Were/weren't leveling in a vernacular English variety. *Language variation and change*, *6*, 273–302.

Senghas, A. (2000). The development of early spatial morphology in Nicaraguan Sign Language. In *Proceedings of the 24th Annual Boston University Conference on Language Development* (pp. 696–707).

Senghas, A., & Coppola, M. (2001). Children creating language: How Nicaraguan Sign Language acquired a spatial grammar. *Psychological Science*, *12*, 323–328.

Senghas, A., Coppola, M., Newport, E. L., & Supalla, T. (1997). Argument structure in Nicaraguan Sign Language: The emergence of grammatical devices. In *Proceedings of the 21st annual Boston university conference on language development* (pp. 550–561). Cascadilla Press Boston, MA.

Shannon, C. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, *27*, 379–423.

Singleton, J. L., & Newport, E. L. (2004). When learners surpass their models: The acquisition of American Sign Language from inconsistent input. *Cognitive Psychology*, *49*, 370–407.

Skinner, B. (1957). *Verbal behavior.*. Appleton-Century-Crofts.

Smith, J., Durham, M., & Fortune, L. (2007). "Mam, my trousers is fa'in doon!": Community, caregiver, and child in the acquisition of variation in a Scottish dialect. *Language Variation and Change*, *19*, 63–99.

Smith, K., Fehér, O., & Ritt, N. (2014). Eliminating unpredictable linguistic variation through interaction. In *Annual meeting of the Cognitive Science Society*.

Smith, K., Kirby, S., & Brighton, H. (2003). Iterated learning: A framework for the emergence of language. *Artificial life*, *9*, 371–386.

Smith, K., Perfors, A., Fehér, O., Samara, A., Swoboda, K., & Wonnacott, E. (2017). Language learning, language use and the evolution of linguistic variation. *Phil. Trans. R. Soc. B*, *372*, 20160051.

Smith, K., & Wonnacott, E. (2010). Eliminating unpredictable variation through iterated learning. *Cognition*, *116*, 444–9.

Thompson, W. D. (2015). *Transmission, induction and evolution*. Ph.D. thesis University of Edinburgh.

van Trijp, R. (2013). Linguistic assessment criteria for explaining language change: A case study on syncretism in German definite articles. *Language Dynamics and Change*, *3*, 105–132.

Vouloumanos, A. (2008). Fine-grained sensitivity to statistical information in adult word learning. *Cognition*, *107*, 729–42.

Wonnacott, E. (2011). Balancing generalization and lexical conservatism: An artificial language study with child learners. *Journal of Memory and Language*, *65*, 1–14.

Wonnacott, E., & Newport, E. (2005). Novelty and Regularization: The Effect of Novel Instances on Rule Formation. In A. Brugos, M. R. Clark-Cotton, & S. Ha (Eds.), {*BUCLD 29*}*: Proceedings of the 29th Annual Boston University Conference on Language Development*. Boston, MA: Cascadilla Press.